

**Brooklyn College**  
**Department of Computer and Information Sciences**

**CISC 4335 [55] Parallel and Distributed Computing**

3 hours; 3 credits

Development of parallel and distributed applications. Examples taken from scientific programming, artificial intelligence, operations research, and management science. Symmetric models and asymmetric models. Client-server. Threaded programming. Data parallelism versus algorithmic parallelism. Issues of resource utilization, program decomposition, performance, probability, debugging. Use of formal methods, fault-tolerance, security and other distributed systems issues.

**Textbooks:**

1. Introduction to CUDA Programming, Volume I, CUDA Technical Training, Nvidia Corp., 2008
2. Parallel Processing with CUDA, Tom R. Halfhill, Microprocessor Report, January 28, 2008.

**Syllabus:**

Week 1 - Introduce ideas of grand challenge problems, parallel computing and GPU architecture. Start reading 2.

Week 2 - Review of thread programming in C, C++ and Java, Handouts

Week 3 - Introduction to the architecture of the Nvidia GPU, read CUDA, Supercomputing for the Masses: Part 1, Rob Farber, April, 2008, <http://www.drdobbs.com/hpc-high-performance-computing/207200659>

Week 4 - Initial Cuda constructs in C programs, discussed the incrementArray.cu program described in CUDA, Supercomputing for the Masses: Part 2, Ron Farber, April, 2008, <http://www.drdobbs.com/cpp/207402986> -

Week 5 - Error handling and global memory, CUDA, Supercomputing for the Masses: Part 3, Ron Farber, May, 2008, <http://www.drdobbs.com/hpc-high-performance-computing/207603131>

Week 6 - Shared memory and synchronization, CUDA, Supercomputing for the Masses: Part 4, Ron Farber, June, 2008, <http://www.drdobbs.com/architect/208401741>

Week 7 - Relationships between local memory and global memory, CUDA, SFTM: Part 5, Ron Farber, June, 2008, <http://www.drdobbs.com/hpc-high-performance-computing/208801731>

Week 8 - Global Memory contention, using the CUDA profiler, CUDA, SFTM: Part 6, July 2008, <http://www.drdobbs.com/architect/209601096>

Week 9 - double precision with the Tesla GPU, CUDA, SFTM: Part 7, August 2008, <http://www.drdobbs.com/hpc-high-performance-computing/210102115>

Week 10 - texture memory on the Nvidia Tesla, CUDA, SFTM: Part 11, March 2009, <http://www.drdobbs.com/hpc-high-performance-computing/215900921>

Week 11 - Using libraries with cuda, CUDA, SFTM: Part 10, January, 2010 <http://www.drdobbs.com/architect/212903437>

Weeks 12 - 14, Each student in the class discussed the code that they were converting to use the Nvidia Tesla Gpus at the HPC center at CSI.